

SAM (REX/Ops) Operations Meeting 2008.09.02

Adam (typist), Krzystof, SteveW, Angela, Robert, more...

Meeting time: 52 minutes

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Where did the summer go?

Releases:

No new releases pending. SAM development roadmap is empty.

Off-road new station release up-coming.

DØ:

DØ Down day is today - but appears to be no down activity.

Events worth noting:

No SAM shifter from Aug. 29 to Aug. 31 (and most of Sept. 1).

Friday (Aug. 29): SAM file deliveries were very slow. Robert found that the selection algorithm ignored empty cache disks '

and he has since improved this (DZSAM-231).

Andrew made a change some time ago (6 mo ago) that made the station less apt to fill empty disks. Went too far and made the station extremely biased against empty disks. With the move of cache servers from FCC to GCC and some dying and being restored (lost 14 TB on each cabsrv1 and 2). This made everything slow. Robert rewrote the disk selection algorithm on Friday (should now be slightly biased for empty disks) with a weighting scheme (the old algorithm could get unstable). Robert installed it on Friday and tweaked it over the weekend to fix weights. Now load has shifted on enstore, where it belongs -- and now load is spread on empty machines. [Right now sam-lto2 is not overloaded! - good]

Monday (Sept. 1): "Another meltdown" to SAM as M. Diesburg called it ... after Robert's efforts, there are already good

signs of life at the time of writing (DZSAM-234).

Multiple independent problems - samgfw02 spontaneously rebooted (before reboot lots of errors and segfault from Globus gtk). Services did not automatically restart. Robert started. Operating fine now.

Problem #2: SAM station was not delivering files (but was responsive). Happened because srm stager for srm services at Purdue was dead and that hung the station [station was waiting forever for a reply]. So the station main work queue was blocked. Robert killed the srm stager and that released the station (this will restart the stager). [Bug in SRM stager?]. Stager was yellow in SAAG. Killing the stager destroyed the evidence - but Andrew should look anyway. Purdue SRM was very sick yesterday. Purdue has an SRM v2 dCache. No Condor problems. SRM at UNL was also broken, and SRM at UCSD was sick.

Art: dCache has a program called srmWatch.

Recent unresolved issues:

DZSAM-235: "Lancaster-LCG not visible on SAMGRID monitoring page", a new issue/complaint.

Steve will look into this. Some reports that jobs on LCG are magically working.

DZSAM-234: "Another meltdown" --- has been actively worked on today.

DZSAM-231: Herb complained about file undelivery problem. Robert put in a new algorithm... Is this issue considered as resolved ?

DZSAM-225 : "d0srv071 does not deliver files" -- related to the unresolved DZSAM-54

DZSAM-224 : "samgrid.fnal.gov having serious problems" -- similar to DZSAM-234 ??
Filesystem problems. Follow up with Steve Timm.

CDF:

CDF SAM/DH Shift Report 08/26/08 - 09/01/08

dCache fcdldata147 and fcdldata320 were offline and enstore mover LTO3_10 was down at the beginning of the shift.

There was an attempt to bring fcdldata147 up on Thursday, but the machine failed again. LTO3_10 is back in service since Thursday.

dCache fcdldata316 and fcdldata323 were moved from FCC to GCC and were offline between Wednesday and Monday.

dCache fcdldata109 went offline since about 0010 on Sunday; it is unpingable and noted on FEF's faultlog. Large ~120 difference between requested and queued restores likely due to fcdldata109 being down.

Large restore backlog Tuesday-Friday and again since Saturday.

Some slowness of the db (or a network problem) at about 8am on Sunday visible in the db/dbserver plots, plus a time correlated station message (and a gap in the station plot):

08/31/08 08:04:50 cdf-caf.SM.FileUncachedInDBTimer 6353: Cannot uncache file in DB: CORBA::UserException caught in FileUncachedInDBTimer::dolt()

CSL To Enstore Status marked as DEAD; no transfers on the plots since ~2200 on Sunday; (experts are aware of the situation)

Angela is checking dead monitoring? - Angela can't log into machine - it's is pingable though. Opening a helpdesk ticket.

Shifters: Krzysztof (and John on Friday)

Ray Culbertson's production problem. 8/31 at 0721 exceeded max # of sessions. Also happened at 9/01 2352. CDF Metering code was blocking access to sam_read account. DB limit is 665. Why is the farm using sam_read? Art: There is a environment variable called USE_SAM_METADATA (if TRUE then this is BAD BAD BAD - uses old code and goes in through sam_read account). This variable should not be set.

Krzysztof: There is a tool called pick_event that would also do this.

MINOS:

General Collaboration meeting this week (Thu-Sun) at Fermilab

Enstore -

DCache -

SAM -

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 2071: quadratic project startup time growth

PLAN -

No news!

DB:

See CDF above under Ray.

Got the sam_read # session exceeded error (BUT THIS MAY BE MISLEADING - SEE BELOW). But also got trigger errors.

Program noted by Oracle is /fbsng/condor_local/execute/dir_24518/work/bin (just the local directory for the executable). /local/stage1/condor/dir_nnnnn/work/bin .

Can see the queries. Nelly will look. Performance on cdfopr2 is slow because of this.

An example query is...

```
SELECT /*+ RULE; */ SUM(LUM_INTEGRAL_ONLINE) ,  
SUM(LUM_INTEGRAL_OFFLINE) FROM DATA_FILES JOIN  
DATA_FILES_LUMBLOCKS USING(FILE_ID) JOIN  
FILECATALOG.CDF2_RUNSECTIONS ON( ID BETWEEN LUM_MIN  
AND LUM_MAX) WHERE FILE_NAME = 'er040143.02eaphys'
```

Looks like a production query.

Also a trigger error for sam_prod_ap_update (??). - Maybe this was a connection problem.

User I/O on the machine is high, generating some waiting.

Right now waitstatus of 6 and idle 25%; normal is waitstatus of 2, idle of 40%

No communication with Ray yet.

**** BUT - Max sessions exceeded error may be misleading because it happens on commit when you already have a session. This may be a symptom of some other Oracle problem and Oracle just spits out this error. But there are still a lot of samread sessions.

Enstore/dCache:

No report

Special: Power outage on 9/18 at GCC

- Thursday, September 18th -- GCC Power Outage
- 480V repair, testing, maintenance - **no generators**
- CRA 1-3 hours, CRB 2-8 hours (CRA = computer room A)
- will bring cooling switch boards back first
- expect GCC networking & tape libraries to remain up on UPS
- GCC/CRC and LCC 480 maintenance tbd (Sept. 17, 18, 19)
- try to schedule LCC 13kV maintenance with LCC 480V maintenance
- date still tentative based on availability of FESS resources

